

Toward a More Practical Marking Scheme for IP Traceback

Chao Gong and Kamil Sarac

Department of Computer Science, University of Texas at Dallas, USA

{cxg010700, ksarac}@utdallas.edu

Abstract

Probabilistic packet marking (PPM) has been studied as a promising approach to realize IP traceback. In this paper, we propose a new PPM approach that improves the current state of the art in two practical directions: (1) it improves the efficiency and accuracy of IP traceback and (2) it provides incentives for ISPs to deploy IP traceback in their networks. Our PPM approach employs a new IP header encoding scheme to store the whole identification information of a router into a single packet. This eliminates the computation overhead and false positives due to router identification fragmentation. Our approach does not disclose the IP addresses of the routers having marked packets, thereby alleviating the ISP's security concern of disclosing network topology. Our approach is able to control the distribution of marking information. So it is suitable to be implemented as a value-added service which may create revenue for ISPs. Therefore our PPM approach improves the performance and practicability of IP traceback.

1 Introduction

IP traceback is a technique for tracing the paths of IP datagrams back toward their origins. IP traceback is not a goal but a means to defending against denial-of-service (DoS) attacks. Identifying the origins of attack packets is the first step in making attackers accountable. In addition, after figuring out the network path which the attack traffic follows, the victim under DoS attack can apply defense measures such as packet filtering further from the victim and closer to the source. That improves the efficacy of defense measures and reduces the collateral damage to innocent traffic.

Many IP traceback techniques have been proposed [1, 2, 3]. Among them, the probabilistic packet marking (PPM) approach has been studied mostly [3, 4, 5, 6, 7, 8]. In a PPM approach, the router probabilistically marks packets with its identification information, and then the whole network path is reconstructed by combining a modest number of marked packets.

There are two problems hindering the deployment of PPM approaches in the Internet. First, current PPM approaches have limited efficiency and accuracy in tracing large-scale distributed DoS (DDoS) attacks. Because of the limited marking space in IP header, the router splits its identification information into multiple fragments and marks the packet with one of those fragments. Those fragments need to be reassembled at the destination to restore the router identification. In a DDoS attack, the attack traffic originates from multiple sources and the victim receives identification fragments from multiple routers in different attack paths. The victim needs to verify the correctness of all combinations of fragments and reconstruct attack paths based on correct fragment combinations. The process of combining router identification fragments and verifying their correctness incurs computation overhead on victims and false positives in reconstructed attack paths.

The second problem is that Internet service providers (ISPs) lack incentives to deploy PPM approaches in their networks. Although some end users may clamor for IP traceback for DoS defense, ISPs are reluctant to support PPM if they cannot sell PPM-based IP traceback as a service. Moreover, similar to topology probing tools such as *traceroute*, the PPM approach reveals the network topology information of the ISP supporting it. ISPs generally regard their network topologies as confidential. So a cautious ISP would not like to support PPM, just like it blocks *traceroute* requests.

In this paper, we present an accurate and secure PPM (ASPPM) approach that addresses the above-mentioned problems. ASPPM identifies routers with assigned ID numbers instead of IP addresses. ASPPM also employs a new IP header encoding scheme to store the complete router identification information into a single packet. In ASPPM, if a marked packet is to be forwarded to a customer not purchasing IP traceback service, the marking information in the packet will be removed. Hence, ASPPM is suitable to be deployed by ISPs as a value-added service.

Recording the complete router identification information within a single packet eliminates the computation overhead and false positives which result from combining router identification fragments. This also reduces the number of pack-

ets required for reconstructing attack paths since a router in attack path can be derived from just one marked packet instead of multiple ones. Representing routers with assigned ID numbers instead of IP addresses preserves the confidentiality of ISP networks. Providing IP traceback to users as a value-added service creates revenue for ISPs. Our solution addresses the needs of both end users and ISPs.

The rest of this paper is organized as follows. We review PPM approaches in Section 2. Section 3 describes the ASPPM approach in detail. Section 4 evaluates the performance of ASPPM through mathematical analysis and simulation. Section 5 discusses the issues involved in deploying ASPPM. Section 6 surveys related work. Finally we conclude our work in Section 7.

2 Probabilistic Packet Marking

Burch *et al.* [9] suggested the possibility of IP traceback based on packet marking. The intuition is to notify the packet destination of the network path by recording the existence of the routers on the route in forwarded packets. One feasible packet marking scheme is that the router probabilistically marks packets with its identification information as they are forwarded by that router. The marking information overloads a rarely used field in IP header. While each marked packet represents only a small portion of the path it has traversed, the whole network path can be reconstructed by combining a modest number of marked packets. This kind of approach is referred to as probabilistic packet marking (PPM) [3].

Because of the probabilistic nature of PPM, a packet may arrive at the destination without having been marked by any of the intermediate routers. Wily attackers are able to “insert” false routers into the network path by sending packets with carefully forged marking values. Most PPM approaches reserve a distance field in the marking space to limit the effect of fake marking values. When a router decides to mark a packet, it writes a zero into the distance field; otherwise, the router increments the distance field using a saturating addition. In this way, any packet written by the attacker will have a distance greater than the length of the true attack path. Therefore, it is impossible for an attacker to forge a router closer than the first traceback enabled router through which its packets have to pass.

In a DDoS attack, there are multiple attackers and the attack traffic traverses multiple paths before converging at the victim. The goal of IP traceback is to reconstruct the *attack tree* which is rooted at the victim and composed of the attack paths from all of the attackers to the victim. Therefore, in order to track multiple attackers in a DDoS attack, the PPM approach needs a mechanism to classify the routers in different attack paths. Two kinds of schemes are employed in PPM approaches to reconstruct attack trees. One

is edge marking, the other is node marking supplemented with a network map. In the edge marking scheme, which is used in CEFS [3] and AMS [4], a marked packet carries the information about an edge in the network path. An edge is represented with the two routers at each end of a link. This scheme can distinguish multiple attack paths because the edges in the same path can be jointed together and the routers in different paths produce disjoint edges. In the node marking scheme, which is used in FIT [6], a marked packet carries only the information of an individual router. The victim consults its upstream router map (a tree topology rooted at the victim) to discern routers in different paths.

The advantages of PPM approaches include low overhead on routers, no additional network traffic, and supporting incremental deployment. However, there are two challenges in applying PPM approaches for IP traceback in practice.

(1) Scalability. Current PPM approaches are not scalable to large-scale DDoS attacks. There is no place in current IP header designated to store marking information. To store marking information in an IP option is not feasible because most routers handle packets with IP options very slowly. In PPM approaches, the marking information overloads a rarely used field in IP header, i.e., 16-bit IP identification field. A single packet usually cannot fit the identification information of a router (e.g., a 32-bit IP address or an IP address hash with similar length). The usual solution is to split the router identification into multiple non-overlapping fragments. When a router decides to mark a packet, the router randomly selects one fragment and marks the packet with the selected fragment plus its offset in the original identification. Those fragments are reassembled at the receiver to restore the router identification. In a DDoS attack, the attack traffic originates from multiple sources and the victim receives identification fragments from multiple routers at the same distance. The victim needs to try all combinations of the fragments at each distance with disjoint offset values, check their correctness, and then accepts correct ones.

There are two kinds of schemes to verify the correctness of fragment combinations. One scheme is to include an integrity verification code in marking values, such as a hash [3] or a checksum [5] of router identification. All packets marked by the same router carry the integrity verification codes which are identical or compatible with each other. The other scheme is to use candidate router sets derived from marking values [4, 6]. In this scheme, the victim maintains an upstream router map and all routers at the same distance from the victim in the map are grouped together as a candidate router set. If a combination of the router identification fragments at a distance corresponds to one element in the candidate router set of the same distance, it is accepted as valid.

Neither scheme is 100% accurate, more or less, in veri-

ifying the correctness of fragment combinations [3, 4, 5, 6]. False positive fragment combinations introduce nonexistent routers in reconstructed attack paths. In addition, the process of combining router identification fragments and verifying their correctness incurs computation overhead on the victim. The more the attackers in a DDoS attack, the higher the computation overhead and the more the number of false positives. Hence, router identification fragmentation prevents PPM approaches from being scalable to large-scale DDoS attacks.

(2) Incentives. ISPs lack incentives to deploy PPM approaches in their networks. In general, ISPs are not willing to support a new protocol that cannot be sold as a service. IP traceback accelerates victim’s reaction to DoS attacks and improves the efficacy of DoS defense measures. Although some customers may clamor for IP traceback, it is not easy for ISPs to offer PPM-based IP traceback as a value-added service to create benefit. Since it is unrealistic to maintain per-flow state at routers, the routers supporting PPM have to mark each forwarded packet with same probability, disregarding whether the packet destination is paying for IP traceback service. ISPs need a mechanism to restrict the use of IP traceback service only to paying customers.

More importantly, ISPs would not like to disclose the details of their networks because of security concerns. In current PPM approaches, the router marks packets with its IP address or related variants (e.g., hash of IP address). Any dedicated end system can construct an upstream router map and derive the IP addresses of those routers in the map based on the marking information in received packets. Attackers may utilize that mapping feature to set ISP’s routers as targets.

3 ASPPM Approach

The accurate and secure PPM (ASPPM) approach is a node marking scheme. The end system maintains an upstream router map using the packet marking information collected during peace time. Same as previous PPM approaches, ASPPM reserves a distance field in the marking space to deal with fake marking values. Although ASPPM is similar to FIT [6] in some aspects, ASPPM employs a different marking algorithm and deployment scheme to improve IP traceback performance and appeal to ISPs.

First, ASPPM does no longer identify routers by their IP addresses or related variants. Instead, the identification of a router consists of two parts, its autonomous system (AS) number and an ID number uniquely assigned within the AS. Hence, an end system cannot infer router IP addresses from the marking values in received packets. That alleviates the ISP’s security concern of disclosing network topology.

Second, ASPPM stores the entire identification information of a router in a single packet. Therefore, the router does

not need to split its identification information into multiple fragments. This eliminates the computation overhead and false positives due to router identification fragmentation. It also reduces the number of packets required for reconstructing attack paths since only one marked packet is enough to identify a router in attack path.

Third, our marking algorithm facilitates distinguishing marked packets and the unmarked packets originated from legitimate users. That feature helps to filter out noise during map construction process and improves the accuracy of upstream router map.

Last, ASPPM can restrict packet marking information to interested customers only. A network service which can be restricted to a subset of customers is suitable to be offered as a value-added service. So IP traceback based on ASPPM can create revenue for ISPs as a value-added service.

3.1 Packet Marking

In PPM approaches, the routers supporting PPM mark packets and reconstructed attack paths are composed of only those *traceback enabled* routers. So we argue that it is enough for PPM approaches to identify only traceback enabled routers instead of all routers in the Internet. In ASPPM, each traceback enabled router has a nonzero 13-bit ID number which is uniquely assigned within its AS. Nonzero 13-bit numbers are sufficient to represent 8191 routers. Since only traceback enabled routers need ID numbers, 13 bits are enough for any AS provided that there are no more than 8191 traceback enabled routers in the AS. An AS is identified by its global unique 16-bit AS number. So a traceback enabled router is identified by its AS number and its ID number within the AS.

We borrow the one bit distance scheme from FIT [6]. One bit in the marking space is reserved as the distance field. When a router decides to mark a packet, the router replaces the five least significant bits of the TTL field in the packet with a global constant c , and copies the sixth least significant bit of the TTL field to the 1-bit distance field of the packet. The TTL field is decremented by one each time the packet is forwarded by a router. The packet destination can calculate the hop counts from the router which marked the packet based on the TTL value and distance field value in the packet as well as the constant c . For every forwarded packet, the router calculates a *marking predicate* based on the TTL value, distance field value, and constant c . The marking predicate will be true if the packet has not been marked for the past 32 hops. If the marking predicate is true, the router will mark the packet deterministically. Otherwise, the router will mark the packet probabilistically. The purpose of the marking predicate, similar to the saturating addition on the distance field in other PPM approaches, is to prevent attackers from forging a router closer than the first

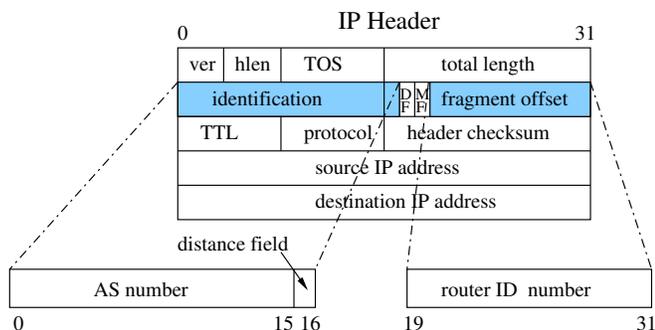


Figure 1. Encoding marking information into IP header.

traceback enabled router in the attack path. We refer readers to [6] for details.

In ASPPM, the marking value consists of 29-bit router identification (16-bit AS number and 13-bit router ID number) and 1-bit distance value. Figure 1 depicts how the marking information fits into IP header. The AS number overloads the IP identification field, the distance bit overloads the unused flag bit next to the IP identification field, and the router ID number overloads the fragment offset field.

Same as all other PPM approaches, ASPPM reuses the IP identification field which is designated for IP fragmentation. Measurement studies show that less than 0.25% of packets are fragmented in the Internet [10]. Ref. [3] discusses the backward compatibility issues of reusing the IP identification field. ASPPM also reuses the reserved flag bit for packet marking, as suggested in [11]. Because the fragment offset field becomes meaningless when the IP identification field is used for the purpose other than IP fragmentation, it has been proposed to reuse the fragment offset field for packet marking [12, 13]. Since the destination regards the IP datagram with a nonzero value in the fragment offset field as an IP fragment, reusing the fragment offset field may cause the destination host to confuse marked packets with IP fragments. Additional mechanisms are required to avoid this kind of confusion. In ASPPM, the marking information in the fragment offset field of a marked packet will be removed before the packet arrives at the destination. For an end user interested in IP traceback, a specific server operated by the user itself or its ISP intercepts and processes the marking information. We will discuss this issue thoroughly in Section 3.2.

The router marks packets with a global certain probability q . When a router decides to mark a packet, besides storing its identification information in the packet and setting the distance field, the router also sets Don't Fragment (DF) flag to be 1 and More Fragments (MF) flag to be 0. Figure 2 describes the marking algorithm. We omit the details

```

For each packet  $p$ 
  let  $r$  be a random number from  $[0,1)$ 
  calculate the marking predicate  $Q$ 
  IF  $r < q$  OR  $Q = \text{true}$  THEN
     $p.\text{identification} := \text{AS number}$ 
     $p.\text{fragment\_offset} := \text{router ID number}$ 
     $p.DF := 1$ 
     $p.MF := 0$ 
    set the distance bit and TTL field accordingly
  
```

Figure 2. Packet marking algorithm at traceback enabled routers.

of the distance manipulation procedure which is identical with FIT.

The purpose of setting the DF and MF flags is two fold. First, setting the DF flag prevents the marked packet from being fragmented downstream from the marking router. So the marking information will not be removed. Second, setting the DF and MF flags helps the packet destination to distinguish marked packets and the unmarked packets originated from legitimate users. In ASPPM, a packet is regarded as a marked packet if its DF flag is 1, MF flag is 0, and fragment offset field has a nonzero value. This function reduces the false positives in the upstream router map. We will describe the details in Section 3.3.

3.2 Deployment

ASPPM is deployed within individual ASes. Some or all routers of an AS are enhanced to support ASPPM marking procedure, and all edge routers are in charge of the distribution of packet marking information. The marking information can only reach the customers paying for IP traceback service. If edge routers forward a marked packet to a customer network or a local host which does not purchase IP traceback service, edge routers will reset the identification field and fragment offset field to remove the marking information in the packet.

Figure 3 describes the forwarding procedure at the edge routers of ASes deploying ASPPM. Suppose an AS, say A , supports ASPPM and one of its edge routers, r , forwards a marked packet p . If packet p is destined to a local host which is not paying for IP traceback service, the marking information in packet p will be removed; if p is destined to a local host which is paying for IP traceback service, the marking information will remain in the packet. If packet p is being forwarded to a customer network of A , the marking information in packet p will be removed or unchanged, depending on whether the customer network purchases IP traceback service. If packet p is being forwarded to a peer

```

For each marked packet  $p$ 
  let  $N$  be the next hop network
  let  $D$  be the destination network of  $p$ 
  let  $d$  be the destination host of  $p$ 
  IF  $D$  is the current AS THEN
    IF  $d$  purchases IP traceback service THEN
      forward  $p$ 
    ELSE
      remove marking information
      forward  $p$ 
  IF  $N$  is a customer network THEN
    IF  $N$  purchases IP traceback service THEN
      forward  $p$ 
    ELSE
      remove marking information
      forward  $p$ 
  IF  $N$  is a peer or provider network THEN
    IF  $D$  supports ASPPM THEN
      forward  $p$ 
    ELSE
      remove marking information
      forward  $p$ 

```

Figure 3. Packet forwarding algorithm at edge routers of AS deploying ASPPM.

or provider network, router r will check whether the destination network of p supports ASPPM. If the destination network does not support ASPPM, the marking information will be removed. Otherwise, the marking information will remain.

We propose to utilize BGP protocol as the vehicle to distribute the ASPPM deployment information in AS level. So ASes will know the existence of remote ASes supporting ASPPM. Each AS deploying ASPPM advertises its support to ASPPM in a BGP attribute in the network route advertisement. Only one bit is enough to represent the ASPPM deployment status. We propose to utilize the community attribute in BGP to carry the ASPPM deployment status bit. The community attribute is 32 bits in length. The community attribute is transitive and optional, which means that if a router does not understand some new community value, it will still forward the value to next hop.

For end users that are interested in IP traceback service, we propose to use *traceback proxy servers* (TPS) to process packet marking information. The TPS intercepts the marking values in packets before they reach the end users. The TPS constructs the upstream router map during peace time, and reconstructs attack paths upon the request of the end user under DoS attack. There are two possible locations for the TPS. One is at ISP side (e.g., be co-located with ISP edge routers), the other is at user side (e.g., reside on end

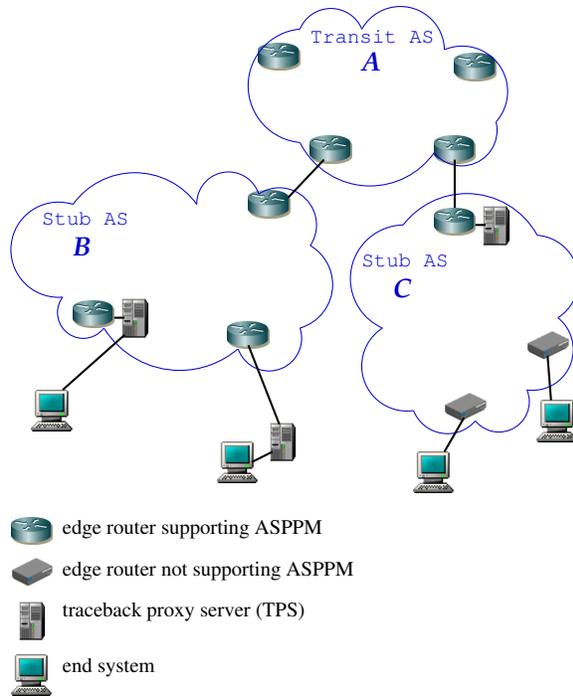


Figure 4. ASPPM deployment. Transit AS A and stub AS B fully deploy ASPPM. All edge routers in ASes of A and B support ASPPM. Stub AS C partially deploys ASPPM. In AS C, only the edge router connected with AS A supports ASPPM.

user firewalls). ISPs may charge more in the former case since ISPs dedicate more resources for serving end users.

If a stub AS has only a few end users interested in IP traceback service, upgrading all its edge routers to support ASPPM may not be desirable. We propose an alternative deployment scheme so that ASPPM can be partially deployed in a stub AS in return for modest performance decrease. The stub AS only needs to upgrade the edge router connected with its provider AS and install a TPS co-located with that edge router. The edge router removes the marking values in all packets except those destined to the end users paying for IP traceback service. The TPS intercepts and processes the marking values in the packets destined to the end users paying for IP traceback. A stub AS partially deploying ASPPM cannot trace DoS attacks originated within its network. Figure 4 illustrates the deployment of ASPPM in AS level.

Each AS deploying ASPPM charges a fee to its customers (networks or end users) who are interested in receiving packet marking information. That only paying customers can get the marking information reduces the attacker's chance to get that information to infer ISP network

#	DF	MF	Offset	Meaning
1	0	0	0	unmarked packet, fragmentation allowed
2	0	0	>0	last fragment
3	0	1	0	first fragment
4	0	1	>0	intermediate fragment
5	1	0	0	unmarked packet, fragmentation not allowed
6	1	0	>0	marked packet
7	1	1	0	improbable
8	1	1	>0	improbable

Table 1. Meaning of combinations.

topologies. It is possible for attackers to manage to get the marking information, for example, through purchasing IP traceback service or perpetrating man-in-the-middle attacks. But attackers cannot derive router IP addresses from packet marking information since the router identification does not reveal IP addresses. Usually, an ISP which is secretive about its topology information is not willing to respond to network topology probing such as *traceroute*. So it is impossible for attackers to map router identifications to IP addresses using those tools.

3.3 Map Construction

Traceback proxy servers (TPS) make use of packet marking information to construct upstream router maps during peace time. The packets in the same TCP connection are viewed as traversing the same network path. Hence, the marking information in those packets is used to construct the network path. This idea was originally proposed in FIT [6]. The marking algorithm employed in ASPPM improves the efficiency and accuracy of map construction procedure.

In FIT, the router is identified by a hash of its IP address and each marked packet carries a fragment of the hash. After collecting enough different identification fragments from a router at a particular distance, the end system needs to scan through the whole IP address space to figure out the IP address of that router. Although some optimizations might be applied to this process (e.g., consider only A, B, and C class IP addresses, or store all IP addresses in a list indexed by their hash values), this mapping process still imposes computation or memory overhead on end systems. Moreover, a packet may traverse the network without being marked by routers. Packet receivers cannot distinguish marked packets and unmarked packets. The end system regards the random values in the marking field of unmarked packets as the marking information from routers and uses them to construct upstream router maps. Those unmarked packets introduce nonexistent routers in constructed upstream router maps.

In ASPPM, routers are represented as ID numbers and one marked packet contains the complete router identifi-

cation. One marked packet unambiguously indicates the router having marked that packet. In addition, a marked packet and an unmarked packet originated from a legitimate user are distinguishable in ASPPM. If a packet has been marked by a router, its DF bit is 1, MF bit is 0, and fragment offset field has a nonzero value. It is unlikely that this setting appears in an unmarked packet originated from legitimate users. An IP datagram with DF flag being 1 will not be fragmented while traversing the network. When the router fragments an IP datagram (its DF flag is 0), the DF flag will not be changed [14]. It is unlikely that hosts send an IP fragment with DF flag set as 1. Table 1 lists the combinations of all possible values of the fragment flags and fragment offset field. Therefore, almost all unmarked packets originated from legitimate users can be distinguished from marked packets. During map construction process, the TPS can filter out most unmarked packets from legitimate users, thereby reducing the false positive routers in upstream router maps.

3.4 Attack Path Reconstruction

When having detected a DoS attack, the victim informs its TPS to reconstruct attack paths based on the marking information in attack packets. The marking value in a single packet reveals the identification of the router having marked that packet and its distance from the victim. Given this information, a router can be pinpointed in the upstream router map easily. Pinpointing a router in the map determines all downstream routers in the network path from that router to the victim, so the process of attack path reconstruction is generally faster in node marking schemes than edge marking schemes. Receiving marked packets from all downstream routers in an attack path makes the victim more confident about the correctness of the path. Requiring fewer packets to reconstruct attack paths means that the victim can react to DoS attacks more quickly.

Depending on the reaction after reconstructing attack paths, the victim may or may not need to contact ASes to map router identifications to IP addresses. If the victim wants to set up packet filters at some remote routers to rate-limit attack traffic, the victim just informs appropriate ASes about the ID numbers of the routers. In this case, there is no need to map router identifications to IP addresses.

4 Evaluation

4.1 Analysis

Given a PPM approach, we assume the marking probability at routers is q and the router identification is divided into k fragments. Suppose in a DDoS attack, the attack traffic traverses m distinct routers at a specific distance d .

In other words, there are m routers which are d hops away from the victim in the attack tree. There will be m^k possible combinations of fragments. Among them, m combinations are correct, and the rest $(m^k - m)$ combinations are incorrect. The computation overhead of combining those fragments and verifying their correctness is bounded by $O(m^k)$. Given a verification scheme, let p be the probability of accepting an arbitrary fragment combination as correct. Then the probability that there are false positive combinations at distance d is:

$$f = 1 - (1 - p)^{m^k - m}.$$

In ASPPM, the complete router identification is stored into a single packet. That is, $k = 1$. Therefore, the computation overhead of combining identification fragments is bounded by $O(m)$ and the probability of false positive combinations $f = 0$. ASPPM does not generate false positive routers in the process of reconstructing attack trees.

4.2 Simulation

Through simulation, we study the number of packets required by ASPPM in (1) upstream router map construction and (2) attack path reconstruction. We simulate that packets traverse network paths of different length and all routers in the network path mark packets with the same probability. We set the marking probability to be 0.04, which is used in most simulation work on previous PPM approaches [3, 4, 6]. We conduct 1000 test runs for each setting of parameters. For comparison purpose, we perform a similar simulation for FIT. We assume that FIT needs 3 marked packets to identify a marking router. That number is the smallest one used in FIT [6].

Figure 5 shows the mean and 95th percentile for the number of packets required to construct a network path in upstream router map. The number of packets required in ASPPM is around one third of that in FIT. Map construction is based on the packets received from the same TCP connection. We cannot always expect to receive enough packets for constructing the entire network path before a TCP connection is closed. Figure 6 depicts the average percentage of the network path being constructed given a certain number of packets. In the simulation, we set the length of network path to be 15 and 25 hops, respectively. Consider a network path which is h hops in length. If the end system constructs the nearest n consecutive routers in the path after receiving a certain number of packets, we say that $\frac{n}{h} \times 100\%$ of the path is constructed. When the number of received packets is small (≤ 200), ASPPM can construct more percentage of network path than FIT.

Figure 7 shows the mean and 95th percentile for the number of packets required to reconstruct an attack path. After identifying a router in an attack path, the victim can

determine all downstream routers in the same path through consulting the upstream router map. Therefore, reconstructing an attack path needs fewer packets compared with constructing a network path of the same length in upstream router map. ASPPM needs around one third as many packets as FIT to reconstruct attack paths. Partially identified attack paths are still valuable for DoS defense. Figure 8 depicts the relationship between the average percentage of the attack path being reconstructed and the number of received packets. We assume the attack path is 15 and 25 hops long, respectively. Suppose an attack path is h hops in length. If the furthest router being identified by the victim after receiving a certain number of packets is m hops away, we say that $\frac{m}{h} \times 100\%$ of the attack path is reconstructed. When receiving a small number of packets (≤ 100), ASPPM can reconstruct more percentage of attack path than FIT.

ASPPM requires fewer packets than FIT for both router map construction and attack path reconstruction. This is because ASPPM needs one marked packet, instead of multiple ones, to determine a router in the network path.

5 Discussion

There are some loose ends that need to be mentioned about our approach. We discuss the most important ones.

5.1 Backward Compatibility

The one bit distance scheme modifies the TTL value in marked packets. The TTL modification may cause a packet to be dropped prematurely before reaching the destination or prevent routers from discarding packets in routing loops. The study in [6] shows that the one bit distance scheme, with a carefully selected TTL replacement constant c , can avoid both of the aforementioned problems.

ASPPM reuses the fragment offset field and the marking value in that field will be removed before the packet reaches the destination. That may make an IP fragment appear to be a non-fragmented IP datagram. ASPPM tries to avoid this mistake by discouraging IP fragment traffic in the network. When an edge router of an AS supporting ASPPM is forwarding a packet, if the packet is an IP fragment (the 2nd, 3rd, 4th, 7th, and 8th lines in Table 1), the router will drop the packet and send back an ICMP “fragmentation needed and the DF bit was set” error message to the source. The source is then expected to reduce the packet size for future packets going toward the same destination to avoid further fragmentation.

5.2 Overhead

If a marked packet is destined to a router, the router should be able to tolerate nonzero values in the fragment

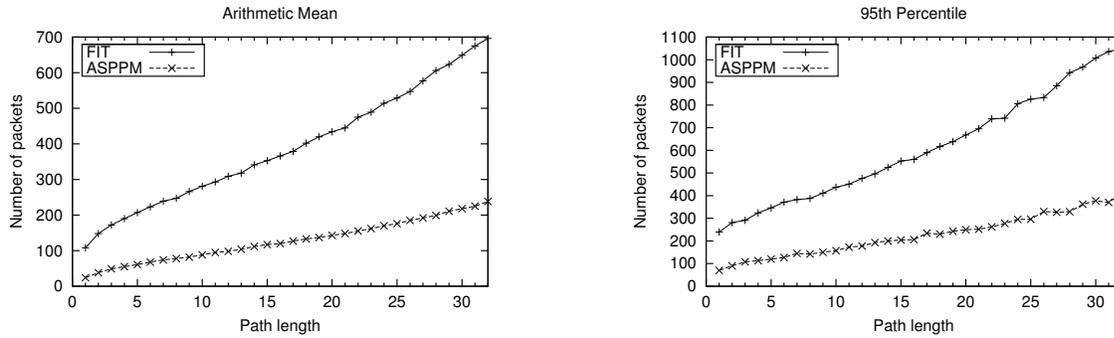


Figure 5. Number of packets required for map construction.

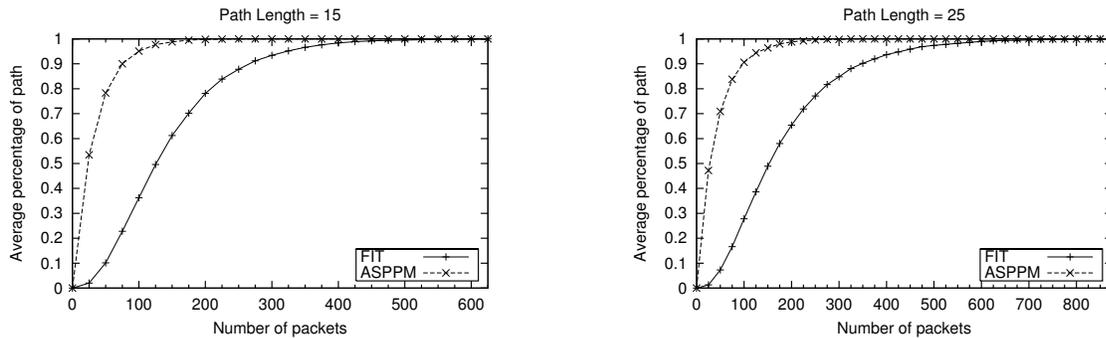


Figure 6. Percentage of map construction.

offset field of the packet. The marking information will not reach the network not supporting ASPPM or the stub AS partially deploying ASPPM. The AS fully deploying ASPPM needs to upgrade all routers to tolerate nonzero values in the fragment offset field. Upgrading all routers in an AS may be difficult in the initial deployment phase. We propose an alternative solution. Suppose an AS begins to deploy ASPPM. The AS upgrades some routers to support ASPPM marking algorithm and modifies all edge routers for controlling the distribution of marking information. The rest legacy routers in the AS network are unchanged. Those traceback enabled routers and the edge routers keep track of the information of legacy routers. The edge routers remove the marking information in the packets destined to the legacy routers. And the traceback enabled routers do not mark the packets destined to the legacy routers. In this way, not all routers in the network need to be upgraded in the initial deployment phase.

There is a close cooperation between the traceback proxy server and end user. In the process of constructing upstream router maps, the traceback proxy server needs to keep track of the TCP connections from and to its users. In the process of reconstructing attack paths, the victim needs to notify the traceback proxy server about the attack signature so that the traceback proxy server can reconstruct attack paths based

on attack packets.

The edge routers of the AS deploying ASPPM control the distribution of packet marking information. That introduces some overhead to edge routers. Because edge routers only keep track of AS-level status information and the status of local hosts, the incurred storage overhead is scalable. Moreover, ASPPM does not introduce additional processing overhead to core routers.

6 Related Work

Motivated by the increasing frequency of DoS attacks and demand for Internet forensic analysis, many IP traceback approaches have been proposed. Based on the location where the network path information is recorded, IP traceback approaches can be classified into three categories:

- packet marking: the information is recorded in forwarded packets.
- ICMP traceback: the information is recorded in router-generated ICMP packets.
- packet logging: the information is recorded at routers.

Burch *et al.* [9] mentioned the idea of IP traceback based on marking packets, either probabilistically or deterministi-

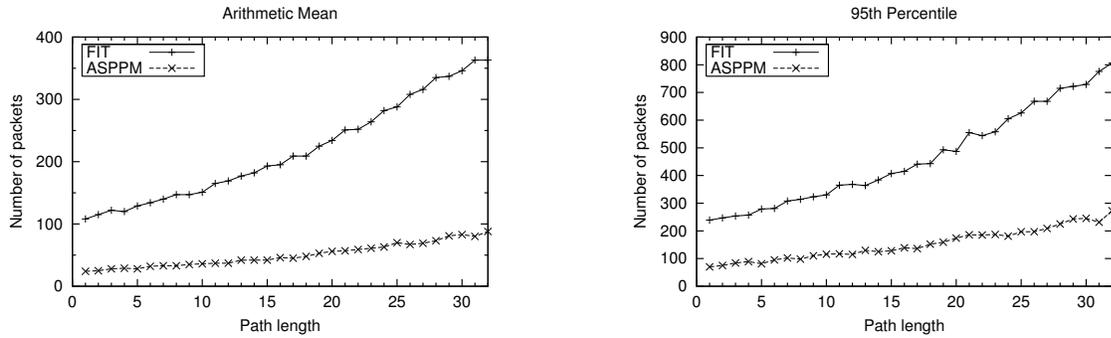


Figure 7. Number of packets required for attack path reconstruction.

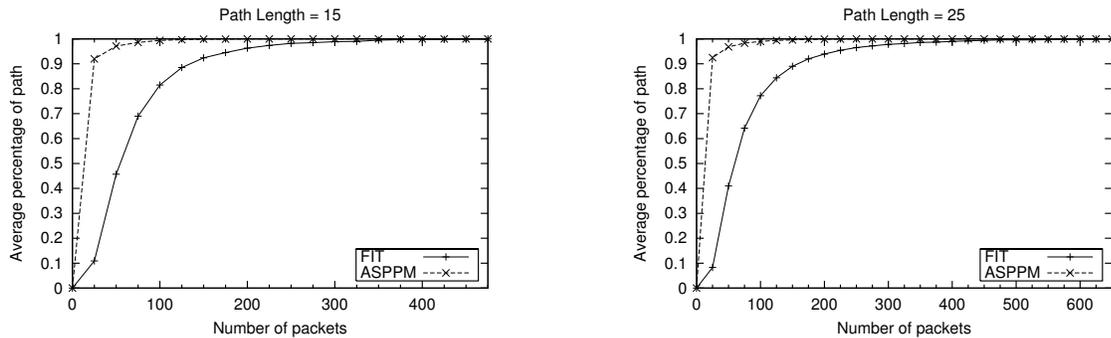


Figure 8. Percentage of attack path reconstruction.

cally, with the identification information (e.g., IP addresses) of the routers they pass through.

Savage *et al.* [3] developed CEFS, a well-known probabilistic packet marking (PPM) approach. CEFS is an edge marking scheme. An edge is encoded as the Exclusive OR (XOR) of the two IP addresses making up the edge. The XOR resulting value is further divided into eight fragments. Each marked packet carries one fragment, the offset of that fragment, and an integrity verification code (a hash value of IP address) which is for reassembling the fragments correctly at the packet destination. Song *et al.* [4] presented AMS to improve CEFS in two aspects. First, a router is identified by the hash value of its IP address. That provides flexibility to encode router identification and to balance the tradeoff between the efficiency and accuracy of attack path reconstruction. Second, AMS uses an upstream router map in place of the integrity verifier to check the correctness of fragment combinations. FIT [6] employs node marking scheme to exploit the full potential of upstream router maps. It also makes use of the TTL field in IP header to calculate the distance of a marking router from the victim so that the distance field in marking space is reduced to one bit.

Some research works have been conducted on the theoretical aspect of PPM approaches. Park *et al.* [7] studied the effectiveness of PPM approaches in the adversarial context

where attackers can forge marking values. Alder [8] theoretically analyzed the tradeoffs between the size of marking space, the number of attackers, and the number of packets required to reconstruct attack paths.

The proposals of IP traceback based on deterministic packet marking (DPM) either abandon the constraint on the marking space [15] or carry extra assumptions (e.g., all ingress routers are traceback enabled) [16]. The major feasible application of DPM is to create a common path signature for all packets traversing the same network route, for the purpose of filtering out attack traffic at the victim site [17, 18, 19].

Bellovin *et al.* [1] proposed ICMP traceback (iTrace). The principle idea is that routers select packets with low probability ($1/20,000$), and then send ICMP traceback messages including the contents of sampled packets and local path information to the same destinations as the selected packets. Mankin *et al.* [20] introduced a feedback mechanism into iTrace. With that feedback mechanism, the victim is able to inform routers of its interest to receive ICMP traceback messages. So routers can create and send ICMP traceback messages purposefully instead of randomly. Barros [21] suggested a simple enhancement to iTrace to address DDoS reflector attacks [22]. The suggestion is to let routers also send ICMP traceback messages to the sources

of the sampled packets.

Sager [23] suggested to realize IP traceback with packet logging. The main idea is to log packets at the routers through which they pass and derive network paths based on the logging information at routers. Historically, packet logging was thought to be impractical because of enormous storage space for packet logs. Snoeren *et al.* [2] presented hash-based IP traceback which reduces the storage overhead significantly. Their approach records packet digests, instead of packets themselves, in a space-efficient data structure, Bloom filter [24]. Some enhancements on hash-based IP traceback have been proposed to further reduce the storage overhead of packet digests [25, 26, 27].

7 Conclusion

In this paper, we have presented a new probabilistic packet marking approach for IP traceback. Our approach stores the whole router identification within a single packet. Therefore, there is no computation overhead and false positives resulting from router identification fragmentation. Our approach does not disclose the IP addresses of the routers having marked packets so as to preserve the confidentiality of ISP networks. Our approach can restrict the marking information to paying customers only. So it is suitable to be implemented as a value-added service to create revenue for ISPs. Our approach improves the performance of IP traceback and provides incentives to ISPs.

One issue in our future work is the pricing of the IP traceback service between ASes. Since the provider AS usually cannot get the information on the expense and earning of its customer ASes, the provider AS can only charge a flat fee for its IP traceback service. An overcharged fee may discourage customer ASes from purchasing the service; and an undercharged fee may dampen the passion of provider AS.

References

- [1] S. Bellovin, M. Leech, and T. Taylor, "ICMP traceback messages." Internet draft, October 2001. <http://search.ietf.org/internet-drafts/draft-ietf-itrace-05.txt>.
- [2] A. Snoeren, C. Partridge, L. Sanchez, C. Jones, F. Tchakountio, B. Schwartz, S. Kent, and W. Strayer, "Single-packet IP traceback," *IEEE/ACM Transactions on Networking*, vol. 10, no. 6, pp. 721–734, 2002.
- [3] S. Savage, D. Wetherall, A. Karlin, and T. Anderson, "Network support for IP traceback," *IEEE/ACM Transactions on Networking*, vol. 9, no. 3, pp. 226–237, 2001.
- [4] D. Song and A. Perrig, "Advanced and authenticated marking schemes for IP traceback," in *Proc. of IEEE INFOCOM*, (Anchorage, AK, USA), April 2001.
- [5] M. Goodrich, "Efficient packet marking for large-scale IP traceback," in *Proc. of ACM Conference on Computer and Communications Security*, (Washington, DC, USA), November 2002.
- [6] A. Yaar, A. Perrig, and D. Song, "FIT: Fast Internet traceback," in *Proc. of IEEE INFOCOM*, (Miami, FL, USA), March 2005.
- [7] K. Park and H. Lee, "On the effectiveness of probabilistic packet marking for IP traceback under denial of service attack," in *Proc. of IEEE INFOCOM*, (Anchorage, AK, USA), April 2001.
- [8] M. Adler, "Trade-offs in probabilistic packet marking for IP traceback," *Journal of the ACM*, vol. 52, no. 2, pp. 217–244, 2005.
- [9] H. Burch and B. Cheswick, "Tracing anonymous packets to their approximate source," in *Proc. of USENIX Systems Administration Conference*, (New Orleans, LA, USA), December 2000.
- [10] I. Stoica and H. Zhang, "Providing guaranteed services without per flow management," in *Proc. of ACM SIGCOMM*, (Cambridge, MA, USA), August 1999.
- [11] D. Dean, M. Franklin, and A. Stubblefield, "An algebraic approach to IP traceback," *ACM Transactions on Information and System Security*, vol. 5, no. 2, pp. 119–137, 2002.
- [12] Z. Gao and N. Ansari, "Enhanced probabilistic packet marking for IP traceback," in *Proc. of IEEE GLOBECOM*, (St. Louis, MO, USA), November 2005.
- [13] M. Muthuprasanna and G. Manimaran, "Space-time encoding scheme for DDoS attack traceback," in *Proc. of IEEE GLOBECOM*, (St. Louis, MO, USA), November 2005.
- [14] "Internet protocol." RFC 791.
- [15] B. Al-Duwairi and T. Daniels, "Topology based packet marking," in *Proc. of International Conference on Computer Communications and Networks*, (Chicago, IL, USA), October 2004.
- [16] A. Belenky and N. Ansari, "IP traceback with deterministic packet marking," *IEEE Communications Letters*, vol. 7, no. 4, pp. 162–164, 2003.
- [17] A. Yaar, A. Perrig, and D. Song, "Pi: A path identification mechanism to defend against DDoS attacks," in *Proc. of IEEE Symposium on Security and Privacy*, (Oakland, CA, USA), May 2003.
- [18] Y. Kim, J. Jo, and F. Merat, "Defeating distributed denial-of-service attack with deterministic bit marking," in *Proc. of IEEE GLOBECOM*, (San Francisco, CA, USA), November 2003.
- [19] Z. Gao, N. Ansari, and K. Anantharam, "A new marking scheme to defend against distributed denial of service attacks," in *Proc. of IEEE GLOBECOM*, (Dallas, TX, USA), November 2004.
- [20] A. Mankin, D. Massey, C. Wu, S. Wu, and L. Zhang, "On design and evaluation of intention-driven ICMP traceback," in *Proc. of International Conference on Computer Communications and Networks*, (Scottsdale, AZ, USA), October 2001.
- [21] C. Barros, "A proposal for ICMP traceback messages." Internet draft, September 2000. <http://www.research.att.com/lists/ietf-itrace/2000/09/msg00044.html>.
- [22] V. Paxson, "An analysis of using reflectors for distributed denial-of-service attacks," *ACM Computer Communication Review*, vol. 31, no. 3, pp. 38–47, 2001.
- [23] G. Sager, "Security fun with OCxmon and cflowd," in *Internet 2 working group meeting*, November 1998.
- [24] B. Bloom, "Space/time trade-offs in hash coding with allowable errors," *Communications of ACM*, vol. 13, no. 7, pp. 422–426, 1970.
- [25] J. Li, M. Sung, J. Xu, L. Li, and Q. Zhao, "Large-scale IP traceback in high-speed Internet: Practical techniques and theoretical foundation," in *Proc. of IEEE Symposium on Security and Privacy*, (Oakland, CA, USA), May 2004.
- [26] T. Lee, W. Wu, and W. Huang, "Scalable packet digesting schemes for IP traceback," in *Proc. of IEEE International Conference on Communications (ICC)*, (Paris, France), June 2004.
- [27] C. Gong and K. Sarac, "IP traceback based on packet marking and logging," in *Proc. of IEEE International Conference on Communications (ICC)*, (Seoul, Korea), May 2005.